Adaptive Submodular Ranking

Fatemeh Navidi* Prabhanjan Kambadur† Viswanath Nagarajan*

June 7, 2016

arXiv:1606.01530v1 [cs.DS] 5 Jun 2016

Abstract

We study a general adaptive ranking problem where an algorithm needs to perform a sequence of actions on a random user, drawn from a known distribution, so as to "satisfy" the user as early as possible. The satisfaction of each user is captured by an individual submodular function, where the user is said to be satisfied when the function value goes above some threshold. We obtain a logarithmic factor approximation algorithm for this adaptive ranking problem, which is the best possible. The adaptive ranking problem has many applications in active learning and ranking: it significantly generalizes previously-studied problems such as optimal decision trees, equivalence class determination, decision region determination and submodular cover. We also present some preliminary experimental results based on our algorithm.

1 Introduction

Many tasks in machine learning can be represented as sequential decision processes. An algorithm is given an *a priori* distribution $\mathcal{D}$ over inputs, and its goal is to satisfy the realized input $i^* \in \mathcal{D}$. In each step, the algorithm chooses an action which partially satisfies i^* and also provides a feedback depending on i^*. This feedback is then used to refine the distribution of i^* for the subsequent steps. So an algorithm in this setting (also called policy) is an adaptive sequence of actions.

Furthermore, many different criteria to satisfy the realized input i^* can be modeled as covering a *submodular* function. Submodular functions are a very general class of set-functions that have certain convexity as well as concavity properties [24]. These functions are used to model utilities in game theory, influence maximization in social networks, diversity in search ranking etc.

In this paper, we study an adaptive optimization problem in the setting described above which simultaneously generalizes many previously-studied problems such as optimal decision trees [20, 23, 10, 8, 18, 9], equivalence class determination [14, 7], decision region determination [22] and submodular ranking [3, 21]. We obtain an algorithm with the best-possible approximation guarantee. Moreover, our algorithm is very simple to state and implement. We also present some preliminary experimental results which are promising.

Before defining the *adaptive submodular ranking* problem formally, we discuss two special cases with applications in active learning and search ranking.

Optimal Decision Trees. This is a basic problem in active learning, which is also called entity identification. There are m possible hypotheses with an *a priori* distribution $\mathcal{D}$ on the true hypothesis i^*. The distribution $\mathcal{D}$ is given by probabilities $\{p_i\}_{i=1}^m$ that sum to one. We can perform tests in order to distinguish between the hypotheses. Each test e costs c_e and is positive for a particular subset T_e of hypotheses. We assume that i^* can be identified by performing all tests.The goal in the optimal decision tree problem is to determine the true hypothesis i^* at the minimum expected cost. Figure 1(a) shows an example with 3 hypotheses, 7 tests and probabilities

*Department of Industrial and Operations Engineering, University of Michigan
†Bloomberg Labs

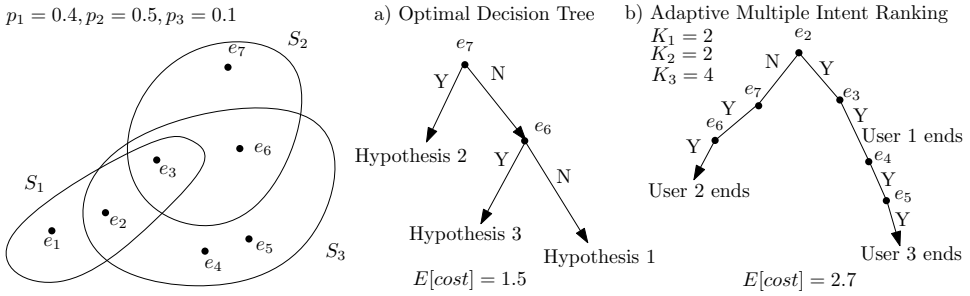

Figure 1: Example for optimal decision tree and adaptive multiple intent ranking problem

$p_1 = 0.4, p_2 = 0.5, p_3 = 0.1$; the sets S_1, S_2, S_3 depict the positive tests for each hypothesis; the decision tree corresponds to a feasible adaptive solution with expected cost 1.5.

Adaptive Multiple Intent Ranking. This is an adaptive version of the search ranking problem studied in [4, 5, 25]. Suppose there are n results to a particular query and m different user-types. Each user-type i is characterized by a subset S_i of the results and a threshold $K_i \leq |S_i|$, which means that a user of type i will be satisfied after seeing at least K_i results from the subset S_i. We are also given a distribution $\mathcal{D}$ on the m user-types. The ranking algorithm displays results one by one, and receives feedback on whether each displayed result e is relevant to user-type i^*, i.e. whether/not $e \in S_{i^*}$. The goal is find an ordering of the n results that minimizes the expected number of results before satisfying a random user $i^* \in \mathcal{D}$. Note that the algorithm needs to balance its effort in learning the identity of i^* and satisfying the requirement of i^*. Figure 1(b) shows an example with $m = 3$, $n = 7$ and probabilities $p_1 = 0.4, p_2 = 0.5, p_3 = 0.1$; the set system depicts the interest-sets S_is and the thresholds are $K_1 = K_2 = 2$, $K_3 = 4$; the decision tree corresponds to a feasible adaptive solution with expected cost 2.7.

For some adaptive optimization problems, one can come up with approximately optimal solutions using static (non-adaptive) solutions that are insensitive to the feedback obtained, eg. [11, 6]. However, this is not the case for the adaptive submodular ranking problem that we consider. Even for the special cases above, there are instances where the optimal adaptive value is much less than the optimal non-adaptive value. So it is important to come up with an adaptive algorithm.

Problem Statement. We start with some basics. A set function $f : 2^U \to \mathbb{R}_+$ on ground set U is said to be submodular if for all subsets $A, B \subseteq U$ we have $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. See [24] for background. The function f is said to be monotone if $f(A) \leq f(B)$ for all $A \subseteq B \subseteq U$. We assume that set functions are given in the standard *value oracle* model, i.e. we can evaluate $f(S)$ for any $S \subseteq U$ in polynomial time.

In the adaptive submodular ranking problem (ASR) we have a ground set U of n elements with costs $\{c_e\}_{e \in U}$. We also have m scenarios. Each scenario $i \in [m] := \{1, \cdots, m\}$ is specified by an interest-set $S_i \subseteq U$ and a normalized monotone submodular function $f_i : 2^{S_i} \to [0, 1]$ where $f_i(\emptyset) = 0$ and $f_i(S_i) = 1$ (every monotone submodular function on S_i can be expressed in this form by scaling and truncation). For notational simplicity, we extend each function $f_i : 2^{S_i} \to [0, 1]$ to arbitrary subsets $S \subseteq U$ by setting $f_i(S) = f_i(S \cap S_i)$. A scenario $i \in [m]$ is said to be covered/satisfied by any subset $S \subseteq U$ of elements such that $f_i(S) = 1$. We are also given a distribution $\mathcal{D}$ on the m scenarios, i.e. probabilities $\{p_i\}_{i=1}^m$ that sum to one: this means that the realized scenario $i^* = i$ with probability p_i, for each $i \in [m]$. The goal in ASR is to find an adaptive ordering of the elements in U that minimizes the expected cost to cover the realized scenario $i^* \in \mathcal{D}$. The key aspect of this problem is that the ranking algorithm receives feedback on whether/not each chosen element e is relevant to user i^* (i.e. whether $e \in S_{i^*}$), and can use this information to choose the next element.

An adaptive solution is represented by a binary decision tree $\mathcal{T}$, where nodes are labeled by elements $e \in U$. Every scenario $i \in [m]$ traces a root-leaf path in the decision tree $\mathcal{T}$ which takes the

yes-branch on any node $e \in S_i$ and the no-branch on any node $e \notin S_i$; let T_i denote the sequence of elements on this path. Every scenario $i \in [m]$ must be satisfied in the decision tree, i.e. $f_i(T_i) \geq 1$. The cost of scenario i in decision tree $\mathcal{T}$ is the cost of the shortest prefix $\bar{T}_i$ of T_i such that $f_i(\bar{T}_i) \geq 1$. The objective in ASR is to minimize the expected cost. We note that multiple scenarios may trace the same path in $\mathcal{T}$.

An important parameter in the analysis of our algorithm is the following:

$$\epsilon \quad := \quad \min_{i \in [m], S \subseteq U} \quad \min_{e \in U : f_i(S \cup e) > f_i(S)} \quad f_i(S \cup e) - f(S). \tag{1}$$

This measures the minimum positive incremental value of any element. This same parameter or its variants appear in all results on the submodular cover problem, eg. [26, 3, 21, 15].

Results. Our main result is an $O(\log \frac{1}{\epsilon} + \log m)$-approximation algorithm for adaptive submodular ranking where $\epsilon > 0$ is as defined in (1) and m is the number of scenarios. Assuming $P \neq NP$, this result is the best possible (up to constant factors) as the set cover problem [12] is a special case of ASR even when $m = 1$. Our algorithm is a simple adaptive greedy-style algorithm. At each step, we assign a score to each remaining element and select the element with maximum score.

Such a simple algorithm was previously unknown even in the special case of optimal decision tree (under arbitrary costs/probabilities), despite a large number of papers [20, 23, 10, 1, 8, 16, 18, 13, 9] on this topic. The first $O(\log m)$-approximation algorithm for ODT was obtained in [18], and this result was extended to the equivalence class determination problem in [9]. Previous results, eg. [23, 10, 1, 8, 16], based on a simple greedy "splitting" algorithm, had a logarithmic dependence on either costs or probabilities which (in the worst case) can be exponential in m.

As direct applications of our algorithm, we obtain approximation algorithms for the two problems mentioned above: an $O(\log(\max_i K_i)+\log m)$-approximation for the adaptive multiple intent ranking problem and an $O(\log m)$-approximation for optimal decision tree. More applications and details can be seen in Section 3.

Related Works. Our work unifies two distinct lines of work in a clean manner. One line of work is on the submodular cover problem and its variants [26, 4, 3, 21]. The other line of work is on the optimal decision tree problem and its variants, eg. [23, 10, 18, 13, 9]. In particular, we combine the time-based analysis for the deterministic submodular ranking problem in [21] with the phase-based analysis for optimal decision tree in [23, 18, 9].

We note that [21] also considers a stochastic variant of submodular ranking, but it is different from ASR because [21] assumes an independent distribution whereas we assume a correlated scenario-based distribution. In particular, unlike ASR, the stochastic submodular ranking problem in [21] does not capture the optimal decision tree problem and its variants.

Recently, [15] also considered a scenario-based submodular cover problem. However, their model requires a *single* submodular function for all scenarios, whereas ASR allows an individual submodular function for each scenario. In this respect our setting is a generalization of [15], eg. ASR captures the submodular ranking problem [3] whereas [15] does not. On the other hand, [15] allows arbitrary feedback whereas ASR as defined only considers binary (yes/no) feedback. We note that our algorithm can be extended to obtain an $O(\log \frac{1}{\epsilon} + \log m)$-approximation in the setting with arbitrary feedback as well – see Section 3.

The notion of "adaptive submodularity" [13] has been very useful in obtaining algorithms for some previously-studied special cases [14, 7, 22] of ASR. In particular, among other results [13] obtained an $O(\log 1/\epsilon + \log 1/p_{min})$-approximation algorithm for ASR when $f_i = f$ for all scenarios i, the function f is adaptive-submodular (a stronger condition than submodular) and $p_{min} = \min_{i=1}^{m} p_i$ is the minimum probability. Our result is stronger in the following ways (i) we allow non-uniform functions f_i for each scenario, (ii) we only require submodularity of these functions, (iii) our performance guarantee $O(\log \frac{1}{\epsilon} + \log m)$ is better. We note that [15] is also an improvement over [13] in points (ii) and (iii) above.

2 The Algorithm

Recall that an instance of ASR consists of a ground set U of elements with costs $\{c_e\}_{e\in U}$, and m scenarios with an interest-set $S_i \subseteq U$, submodular function $f_i : 2^{S_i} \to [0,1]$ and probability p_i associated with each scenario $i \in [m]$. The goal in ASR is to find an adaptive ordering of the elements in U that minimizes the expected cost to cover the realized scenario $i^* \in \mathcal{D}$.

The "state" of our algorithm (i.e. any node in its decision tree) will be represented by (i) the set $E \subseteq U$ of previously displayed elements, and (ii) the set $H \subseteq [m]$ of scenarios that are compatible with feedback (on E) received so far and are still uncovered. At any state (E, H), our algorithm does the following. For each element $e \in U \setminus E$, we define $L_e(H) \subseteq H$ as follows:

$$L_e(H) = \text{argmin}(|\{i \in H : e \in S_i\}|, |\{i \in H : e \notin S_i\}|) \tag{2}$$

Then we select element $e \in U \setminus E$ that maximizes:

$$\frac{1}{c_e} \cdot \left(\sum_{j \in L_e(H)} p_j \; + \; \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} \right). \tag{3}$$

Note that H only contains uncovered scenarios. So $f_i(E) < 1$ for all $i \in H$ and the denominator in the sum above is always positive. Next we analyze the performance of this algorithm. Below, for any subset $T \subseteq [m]$ of scenarios, we use $\Pr(T) = \sum_{i \in T} p_i$.

Let OPT denote an optimal solution to the ASR instance and ALG be the solution found by the above algorithm. Set $L := 15(1 + \ln 1/\epsilon + \log_2 m)$ and its choice will be clear later. We assume (without loss of generality) by scaling that all costs are positive integers. We refer to the total cost incurred at any point in a solution as the *time*. For any $k = 0, 1, \cdots$, we define the following quantities:

- $A_k \subseteq [m]$ is the set of uncovered scenarios of ALG at time $L \cdot 2^k$, and $a_k = \Pr(A_k)$.
- x_k is the probability of uncovered scenarios of OPT at time 2^{k-1}.

Claim 1. *The expected cost of ALG and OPT can be bounded as follows.*

$$C_{ALG} \le L \sum_{k \ge 0} 2^k a_k + L \qquad \text{and} \qquad C_{OPT} \ge \frac{1}{2} \sum_{k \ge 1} 2^{k-1} x_k \tag{4}$$

Proof. In ALG, for all $k \ge 1$, the probability of scenarios for which the cover time is in $(2^{k-1}L, 2^k L]$ is equal to $a_{k-1} - a_k$. So we have:

$$Cov_{ALG} \le \sum_{k \ge 1} 2^k L(a_{k-1} - a_k) = \sum_{k \ge 1} 2^k L a_{k-1} - \sum_{k \ge 1} 2^k L a_k \tag{5}$$

$$= 2 \sum_{k \ge 0} 2^k L a_k - (\sum_{k \ge 0} 2^k L a_k - L) = \sum_{k \ge 0} 2^k L a_k + L \tag{6}$$

Above we use this fact that $a_0 = 1$. On the other hand, in OPT, for all $k \ge 1$, the probability of scenarios for which the cover time is in $(2^{k-2}, 2^{k-1}]$ is equal to $x_{k-1} - x_k$. So we have:

$$Cov_{OPT} \ge \sum_{k \ge 2} 2^{k-2}(x_{k-1} - x_k) = \sum_{k \ge 2} 2^{k-2} x_{k-1} - \sum_{k \ge 2} 2^{k-2} x_k \tag{7}$$

$$= \sum_{k \ge 1} 2^{k-1} x_k - \frac{1}{2}(\sum_{k \ge 1} 2^{k-1} x_k - 1) \ge \frac{1}{2} \sum_{k \ge 1} 2^{k-1} x_k \tag{8}$$

Where we use this fact that $x_1 = 1$. □

Thus, if we could upper bound each a_k by some multiple of x_k, then it is easy to establish the approximation factor. However, this is not the case here and instead we prove:

Lemma 2.1. *For all $k \geq 1$ we have $a_k \leq 0.2a_{k-1} + 3x_k$.*

This lemma implies our main result:

Theorem 2.2. *The algorithm for* ASR *is an $\mathcal{O}(\log 1/\epsilon + \log m)$-approximation algorithm.*

Proof. Let $Q = \sum_{k \geq 0} L \cdot 2^k a_k + L$, which is the bound on C_{ALG} from (4). Using Lemma 2.1:

$$\begin{aligned} Q \;&\leq\; L \cdot \sum_{k \geq 1} 2^k (0.2a_{k-1} + 3x_k) \;+\; L(a_0 + 1) \\ &\leq\; 0.4L \cdot \sum_{k \geq 0} 2^k a_k \;+\; 6L \cdot \sum_{k \geq 1} 2^{k-1} x_k \;+\; L(a_0 + 1) \\ &\leq\; 0.4(Q - L) \;+\; 6L \cdot \sum_{k \geq 1} 2^{k-1} x_k \;+\; 2L \;\leq\; 0.4 \cdot Q + 12L \cdot C_{OPT} + 1.6L \end{aligned} \tag{9}$$

The first inequality in (9) is by definition of Q and $a_0 \leq 1$. The second inequality in (9) uses the bound on C_{OPT} from (4). Finally, as $C_{OPT} \geq 1$ (all costs are positive integers), we have $Q \leq 0.4 \cdot Q + 13.6L \cdot C_{OPT}$. This yields $Q \leq \frac{136}{6} L \cdot C_{OPT}$. Since $L = 15(1 + \ln 1/\epsilon + \log m)$ and $C_{ALG} \leq Q$, we obtain the theorem. □

We now prove Lemma 2.1 for a fixed $k \geq 1$. Consider any time t between $L \cdot 2^{k-1}$ and $L \cdot 2^k$. Note that ALG's decision tree induces a partition of all the uncovered scenarios at time t, where each part H consists of all scenarios that are at a particular node (E, H) at time t. Let $R(t)$ denote the set of parts in this partition. Note that all scenarios in A_k appear in $R(t)$ as these scenarios are uncovered even at time $L \cdot 2^k \geq t$. Similarly, all scenarios in $R(t)$ are also in A_{k-1}.

For any part $H \in R(t)$, let (E, H) denote the node in ALG's decision tree corresponding to H. We note that E consists of all elements that have been completely displayed by time t. The element f selected at this node is *not* included in E (though f starts at/before time t it is not yet completely displayed). Let $T_H(k)$ denote the subtree of OPT that corresponds to paths traced by scenarios in H up to time 2^{k-1}. Note that each node (labeled by element $e \in E$) in $T_H(k)$ has two outgoing branches: one corresponding to $L_e(H)$ and the other to $H \setminus L_e(H)$. We define $\mathsf{Stem}_k(H)$ to be the path in $T_H(k)$ that at each node (labeled e) follows the branch corresponding to $H \setminus L_e(H)$. Below we also use $\mathsf{Stem}_k(H)$ to denote the set of elements that are completely displayed on this path. As all the scenarios in H are compatible with the feedback from elements E,

Observation 1. *Consider any node (E, H) in ALG. Then for all $e \in E$ we have exactly one of the following: $e \in S_i$ for all $i \in H$, or $e \notin S_i$ for all $i \in H$. Hence, $L_e(H) = \emptyset$ for all $e \in E$.*

Definition 1. *Each node (E, H) in ALG is exactly of one of the following types:*

- *"bad" if the probability of uncovered scenarios at the end of $\mathsf{Stem}_k(H)$ is at least $\frac{\Pr(H)}{3}$.*
- *"okay" if it is not bad and the probability of $\cup_{e \in \mathsf{Stem}_k(H)} L_e(H)$ is at least $\frac{\Pr(H)}{3}$.*
- *"good" if it is neither bad nor okay and the probability of scenarios that get covered by $\mathsf{Stem}_k(H)$ is at least $\frac{\Pr(H)}{3}$.*

To see that this is well defined, note by definition of $\mathsf{Stem}_k(H)$ that each scenario in H is (i) uncovered at the end of $\mathsf{Stem}_k(H)$, or (ii) in $L_e(H)$ for some $e \in \mathsf{Stem}_k(H)$, or (iii) covered by some prefix of $\mathsf{Stem}_k(H)$. So the total probability of the scenarios in one of these 3 categories must be at least $\frac{\Pr(H)}{3}$. Therefore each node (E, H) is at least of one of the three types mentioned.

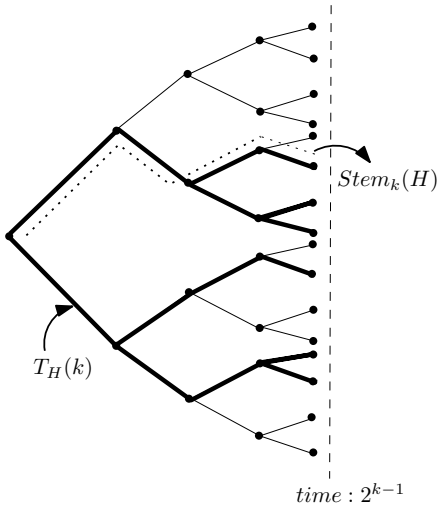

Figure 2: Path $\mathsf{Stem}_k(H)$ in OPT

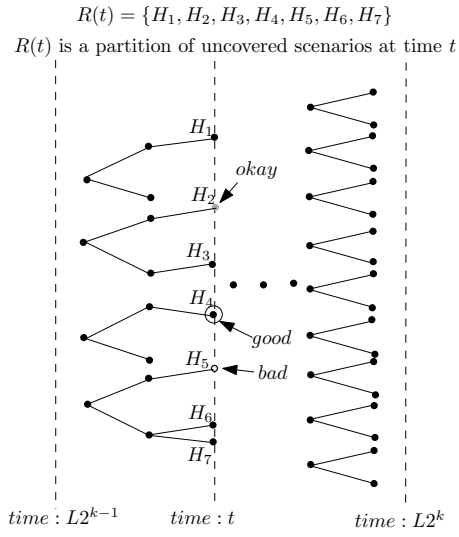

Figure 3: Bad, good and okay nodes in ALG

Observation 2. *For any time $L2^{k-1} < t \leq L2^k$, we have $\sum_{\substack{H \in R(t) \\ H:bad}} \Pr(H) < 3x_k$.*

Proof. Note that $Stem_k(H) \subseteq T_H(k)$. Recall that $T_H(k)$ was the subtree of OPT up to time 2^{k-1} that only contains the scenarios in H. So, the probability of uncovered scenarios at the end of $Stem_k(H)$ is at most the probability of scenarios in H that are not covered in OPT by time 2^{k-1}. This probability is at least equal to $P(H)/3$ based on the definition of bad nodes. Now, since nodes in $R(t)$ denotes a subpartition of scenarios, we have $\sum_{\substack{H \in R(t) \\ H:bad}} \Pr(H)/3 < x_k$. □

The following quantity turns out to be useful in our proof of Lemma 2.1.

$$G := \sum_{t > L2^{k-1}}^{L2^k} \sum_{H \in R(t)} \max_{e \in U \setminus E} \frac{1}{c_e} \cdot \left(\Pr(L_e(H)) + \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} \right) \tag{10}$$

Above, for any $H \in R(t)$ the set E of elements comes from the node (E, H) in ALG corresponding to H. Note that G corresponds to the total "gain" according to our algorithm's selection criterion (3) accrued from time $L2^{k-1}$ to $L2^k$. In what follows we obtain a lower and upper bound for G and combine them to prove lemma 2.1.

Lemma 2.3. *We have $G \geq L \cdot (a_k - 3x_k)/3$*

Proof. Considering only the good/okay nodes in each $R(t)$ in the expression (10) for G,

$$G \geq \sum_{t > L2^{k-1}}^{L2^k} \sum_{\substack{H \in R(t) \\ H:okay}} \max_{e \in U \setminus E} \frac{\Pr(L_e(H))}{c_e} + \sum_{t > L2^{k-1}}^{L2^k} \sum_{\substack{H \in R(t) \\ H:good}} \max_{e \in U \setminus E} \frac{1}{c_e} \cdot \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)}$$

Fix any time t. For any node (E, H) with $H \in R(t)$ define $W(H) = \mathsf{Stem}_k(H) \setminus E$. Recall that the total cost of elements in $\mathsf{Stem}_k(H)$ is at most 2^{k-1}; so $c(W(H)) \leq 2^{k-1}$.

Case 1. (E, H) is an okay node. Since $W(H) \subseteq U \setminus E$ we can write:

$$\max_{e \in U \setminus E} \frac{\Pr(L_e(H))}{c_e} \;\geq\; \max_{e \in W(H)} \frac{\Pr(L_e(H))}{c_e} \;\geq\; \frac{\sum_{e \in W(H)} \Pr(L_e(H))}{c(W(H))} \;\geq\; \frac{\Pr(\cup_{e \in W(H)} L_e(H))}{2^{k-1}}$$
$$= \;\frac{1}{2^{k-1}} \cdot \Pr(\cup_{e \in \mathsf{Stem}_k(H)} L_e(H)) \;\geq\; \frac{\Pr(H)}{3 \cdot 2^{k-1}} \tag{11}$$

The equality in (11) is by Observation 1, and the last inequality is by definition of an okay node.

Case 2. (E, H) is a good node. Below, we use $F \subseteq H$ to denote the set of scenarios that get covered in $\mathsf{Stem}_k(H)$; by definition of a good node, we have $\Pr(F) \geq \Pr(H)/3$. Again using $W(H) \subseteq U \setminus E$, we have:

$$\max_{e \in U \setminus E} \frac{1}{c_e} \cdot \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} \;\geq\; \max_{e \in W(H)} \frac{1}{c_e} \cdot \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)}$$
$$\geq \;\frac{1}{c(W(H))} \sum_{e \in W(H)} \sum_{i \in H} p_i \cdot \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} \;\geq\; \frac{1}{2^{k-1}} \sum_{i \in W(H)} p_i \cdot \sum_{e \in W(H)} \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)}$$
$$\geq \;\frac{1}{2^{k-1}} \sum_{i \in H} p_i \cdot \frac{f_i(W(H) \cup E) - f_i(E)}{1 - f_i(E)} \;=\; \frac{1}{2^{k-1}} \sum_{i \in H} p_i \cdot \frac{f_i(\mathsf{Stem}_k(H)) - f_i(E)}{1 - f_i(E)} \tag{12}$$
$$\geq \;\frac{1}{2^{k-1}} \sum_{i \in F} p_i \cdot \frac{1 - f_i(E)}{1 - f_i(E)} \;=\; \frac{\Pr(F)}{2^{k-1}} \;\geq\; \frac{\Pr(H)}{3 \cdot 2^{k-1}} \tag{13}$$

The inequality in (12) is by submodularity of the f_is, and the equality is by definition of $W(H)$. The first inequality in (13) is by definition of F being the covered scenarios in $\mathsf{Stem}_k(H)$ and the last inequality is by definition of a good node.

Combining lower bounds for okay (11) and good (13) nodes,

$$G \geq \sum_{t > L2^{k-1}}^{L2^k} \sum_{\substack{H \in R(t) \\ H:okay}} \frac{\Pr(H)}{3 \cdot 2^{k-1}} + \sum_{t > L2^{k-1}}^{L2^k} \sum_{\substack{H \in R(t) \\ H:good}} \frac{\Pr(H)}{3 \cdot 2^{k-1}}$$
$$= \sum_{t > L2^{k-1}}^{L2^k} \frac{1}{3 \cdot 2^{k-1}} \left(\Pr(R(t)) - \sum_{\substack{H \in R(t) \\ H:bad}} \Pr(H) \right) \geq \sum_{t > L2^{k-1}}^{L2^k} \frac{a_k - 3x_k}{3 \cdot 2^{k-1}} = \frac{L \cdot (a_k - 3x_k)}{3}$$

The first equality uses the fact that the nodes corresponding to $H \in R(t)$ are exactly one of the types bad/okay/good. The last inequality uses Observation 2 and that $R(t)$ contains all scenarios in A_k. □

Lemma 2.4. *We have $G \leq a_{k-1} \cdot (1 + \ln 1/\epsilon + \log m)$.*

Proof. For any scenario $i \in A_{k-1}$ (i.e. uncovered in ALG by time $L2^{k-1}$) let π_i be the path traced by i in ALG's decision tree, starting from time $2^{k-1}L$ and ending at 2^kL or when i gets covered. For each element e that appears in π_i, let $1 \leq t_{e,i} \leq c_e$ denote the units of time when e is being displayed during the interval $(L2^{k-1}, L2^k]$. Note that there can be at most two elements e in π_i with $t_{i,e} < c_e$: one that is being displayed at time $L2^{k-1}$ and another at $L2^k$.

Recall that every scenario in $R(t)$ appears in A_{k-1}. So only scenarios in A_{k-1} can contribute to

G and we can rewrite (10) as follows:

$$\begin{aligned} G &= \sum_{i \in A_{k-1}} p_i \cdot \sum_{e \in \pi_i} t_{e,i} \cdot \frac{1}{c_e} \left(\frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} + \mathbb{1}[i \in L_e(H)] \right) \\ &\leq \sum_{i \in A_{k-1}} p_i \cdot \left(\sum_{e \in \pi_i} \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} + \sum_{e \in \pi_i} \mathbb{1}[i \in L_e(H)] \right) \end{aligned} \tag{14}$$

Fix any scenario $i \in A_{k-1}$. For the first term, we use Claim 2.1 in [3] which relies on the definition of ϵ in (1). This implies $\sum_{e \in \pi_i} \frac{f_i(e \cup E) - f_i(E)}{1 - f_i(E)} \leq 1 + \ln \frac{1}{\epsilon}$. In order to bound the second term, note that if scenario $i \in L_e(H)$ when ALG selects element e then number of possible scenarios decreases by at least a factor of *two* in path π_i. So such an event can happen at most $\log_2 m$ times along the path π_i. Thus we can write $\sum_{e \in \pi_i} \mathbb{1}[i \in L_e(H)] \leq \log m$. The lemma now follows from (14). □

We now complete the proof of Lemma 2.1. Using Lemma 2.3 and Lemma 2.4 we can write:

$$L \cdot (a_k - 3x_k)/3 \quad \leq \quad G \quad \leq \quad a_{k-1} \cdot (1 + \ln 1/\epsilon + \log m) = \quad a_{k-1} \cdot \frac{L}{15}$$

Rearranging, we obtain $a_k \leq 0.2 \cdot a_{k-1} + 3x_k$ as needed.

3 Extensions

The ASR problem as defined involves binary feedback from each element. Our algorithm can be easily generalized to handle arbitrary (non-binary) feedback as well. As before, we have elements U with costs $\{c_e\}_{e \in U}$ and m scenarios with probabilities $\{p_i\}_{i=1}^m$ (the realized scenario $i^* = i$ with probability p_i). Each element $e \in U$ has an initially unknown state $\delta(e) \in \Gamma$. When an element e is displayed, it provides its state $\delta(e)$ as feedback. Each scenario specifies a state for each element in U, i.e. we are given values $\{d_{i,e} : e \in U, i \in [m]\}$ where $d_{i,e}$ is the state of element e under scenario i. In the binary case considered in Section 2, $|\Gamma| = 2$ (corresponding to yes/no) and each scenario i specifies a "yes" state for elements $e \in S_i$ and a "no" state for elements $e \notin S_i$. Each scenario $i \in [m]$ is also associated with a function f_i defined over subsets of $U \times \Gamma$; we again assume that f_i is monotone submodular and takes values in $[0, 1]$. Note that the ground set of these functions is different from the set U of elements. Scenario $i \in [m]$ is said to be covered by a subset $S \subseteq U$ of elements if $f_i(\{(e, \delta(e)) | e \in S\}) = 1$. The goal is to find an adaptive ordering of the elements so as to minimize the expected cost to cover the realized scenario.

A slight modification of the algorithm in Section 2 also works here. Recall the definition (2) of $L_e(H)$ where $H \subseteq [m]$ is the "current" set of scenarios at any point in the algorithm. We extend this definition as follows. For any element e, let $B_e(H)$ denote the set with maximum cardinality amongst $\{i \in H : d_{i,e} = t\}$ for $t \in \Gamma$; then $L_e(H) = H \setminus B_e(H)$. The algorithm again selects elements according to the criterion (3). The analysis of the $O(\log 1/\epsilon + \log m)$ approximation ratio of this algorithm remains unchanged from Section 2.

4 Applications

In this section we discuss various applications of ASR. Some problems have been studied previously and our result matches the best approximation ratio known with the added advantage of being a simpler (and more efficient) algorithm. Some other problems are new and our result provides the first approximation ratio for these.

Deterministic Submodular Ranking In this problem we are given a set of elements $[n]$ and m monotone submodular functions $f_1, f_2, \ldots, f_m$ such that $f_i : 2^{[n]} \to R_+$ and are normalized between 0 and 1. We also have a weight vector $w \in R_+^m$ and our goal is to find a linear ordering of elements that minimizes $\sum_{i=1}^n w_i c_i$, where c_i is the first time that the value of function f_i reaches one or higher. We can use our algorithm for this problem with the same f_is if we define $S_1 = S_2 = ... = S_m = [n] = U$, and $p_i = w_i/(\sum_{j=1}^n w_j)$. Theorem 2.2 directly gives an $\mathcal{O}(\log m + \log \frac{1}{\epsilon})$-approximation algorithm. Moreover, by observing that in (3) $L_e = \emptyset$ always, we obtain an $\mathcal{O}(\log \frac{1}{\epsilon})$-approximation, which matches the best result known [3].

Adaptive Multiple Intent Ranking This problem was defined in the introduction. It can be modeled as ASR using the functions:

$$f_i(S) = \frac{\min(|S \cap S_i|, K_i)}{K_i}, \quad \text{for all } S \subseteq U \text{ and } i \in [m].$$

Note that each f_i is bounded between 0 and 1, and is monotone submodular. We also have $\epsilon = 1/\max_{i\in[m]} K_i$. So Theorem 2.2 yields an $\mathcal{O}(\log \max_{i\in[m]} K_i + \log m)$-approximation algorithm.

Optimal Decision Tree This problem was also defined in the introduction. It captures many applications in active learning, medical diagnosis, etc. In order to cast this as an instance of ASR, we need some additional definitions. For each test $e \in U$ and hypothesis $i \in [m]$ we define $T_e(i) = \{j \in [m] | e \in S_i \oplus S_j\}$ where $\oplus$ is the symmetric difference operator, and set

$$f_i(S) = |\cup_{e\in S} T_e(i)| \cdot \frac{1}{m-1}, \quad \text{for all } S \subseteq U \text{ and } i \in [m].$$

Note that f_is are coverage functions, so they are monotone and submodular; also the scaling by $\frac{1}{m-1}$ ensures they take values between 0 and 1. By definition of $T_e(i)$, we know that $\bigcup_{e\in S} T_e(i)$ are the hypotheses that differ from i in at least one of the tests S. Hence, having $f_i(S) = 1$ is equivalent to $|\bigcup_{e\in S} T_e(i)| = m - 1$, i.e. we have identified i as the true hypothesis. Here we have $\epsilon = \frac{1}{m-1}$, so Theorem 2.2 gives an $\mathcal{O}(\log m)$-approximation algorithm. This matches the best result known in [18], but their algorithm is more complicated than ours.

Generalized Optimal Decision Tree Our algorithm also extends to the setting when we do not have to identify a unique hypothesis. Here we are given a threshold t_i for each $i \in [m]$ such that it suffices to output a subset of at most t_{i^*} hypotheses that is guaranteed to contain the true hypothesis i^*. This can be handled easily by changing the definition of f_is to:

$$f_i(S) = \min\left\{|\cup_{e\in S} T_e(i)| \cdot \frac{1}{m - t_i}, 1\right\}, \quad \text{for all } S \subseteq U \text{ and } i \in [m].$$

Note that this time we will have $f_i(S) = 1$ if the number of hypotheses that do not match i on tests S is at least $m - t_i$; so this corresponds to having at most t_i possible hypotheses. And Theorem 2.2 implies an $\mathcal{O}(\log m)$-approximation algorithm. To the best of our knowledge this is the first approximation algorithm in this setting; previous results only seem applicable when t_is are uniform.

Equivalence Class Determination We are given an unknown hypothesis in a set of m hypotheses and a partition Q of $[m]$. Let $Q(i)$ be the subset in the partition that contains i. We are allowed to perform a set of tests, and we know that each partition is uniquely identified based on

the results of all tests. Our goal is to minimize the expected cost of tests until we recognize the *part* containing the true hypothesis. We can model this as an ASR instance with

$$f_i(S) = \frac{|\cup_{e\in S}(T_e(i)\cap Q(i)^c)|}{|Q(i)^c|}, \quad \text{for all } S\subseteq U \text{ and } i\in[m].$$

Where A^c denotes the complement of set A. The $T_e(i)$ are as defined above for optimal decision tree. Note that f_is are monotone submodular with values between 0 and 1. Furthermore, $f_i(S)=1$ means that $Q(i)^c \subseteq \cup_{e\in S} T_e(i)$, which means that the set of compatible hypotheses based on the tests S is a subset of $Q(i)$. Again, Theorem 2.2 implies an $\mathcal{O}(\log m)$-approximation algorithm. This is equal to the best previous result [9], but again our algorithm is much simpler.

Decision Region Determination As before, we are given an unknown hypothesis in a set of m hypotheses and tests U. We also have a set of decisions: each decision j is a "region" $D_j \subseteq [m]$ that corresponds to the hypotheses that it is applicable to. The goal is to find a policy for performing tests with minimum expected cost so as to find *any* decision region containing the true hypothesis. For each hypothesis $i\in[m]$ and decision j such that $i\in D_j$ define $f_{i,j}(S) = \frac{|\bigcup_{e\in S}(T_e(i)\cap {D_j}^c)|}{|{D_j}^c|}$. Clearly $f_{i,j}$s are monotone submodular with values between 0 and 1. Also, $f_{i,j}(S)=1$ means that ${D_j}^c \subseteq \bigcup_{e\in S} T_e(i)$, which means that the set of compatible hypotheses based on the tests S is a subset of decision region D_j. However, we may stop when it is determined that the true hypothesis is in any one of the decision regions. This criterion (for hypothesis i) corresponds to at least one $f_{i,j}(S)=1$ among $\{j : i\in D_j\}$ reaches one. In [17] it is shown that the "OR of submodular functions" is submodular. Based on that, we express:

$$f_i(S) = 1 - \prod_{i\in D_j}(1-f_{i,j}(S)), \quad \text{for all } S\subseteq U \text{ and } i\in[m].$$

So we obtain an instance of ASR. Note that here the parameter $\epsilon = \min_i \prod_{j:i\in D_j} \frac{1}{|{D_j}^c|}$ is much smaller due to the OR construction. Still, we have $\epsilon = \Omega(m^{-r})$ where r is the maximum number of decision regions that contain a hypothesis. So in this case, Theorem 2.2 implies an $\mathcal{O}(r\log m)$-approximation algorithm where r is the maximum number of decision regions that contain a hypothesis. [22] obtained an $\mathcal{O}(\min\{r,d\}\cdot \ln\frac{1}{\min_i p_i})$-approximation algorithm for this problem, which was improved by [15] to $\mathcal{O}(\min\{r,d\}\cdot\log m)$; here d is the maximum size of a decision region. Our algorithm runs in time polynomial in m and r, unlike [22, 15] which required time exponential in r. As in [22, 15], we can also obtain an $\mathcal{O}(d\log m)$-approximation with running time exponential in d.

5 Experiments

In this section, we present experimental results for the adaptive multiple intent ranking problem (MIR) and the (generalized) optimal decision tree problem (ODT). We use expected cost (number of elements) as the performance metric to compare different algorithms for the MIR and ODT problems. For example, in the MIR case, if i^* is satisfied after looking at k elements, we say that $cost_{i^*}=k$; the performance metric is then $\sum_{i\in H} p_i\cdot cost_i$.

Environment: We developed high-quality python modules to evaluate the performance of our algorithms against their well-known counterparts for both ODT and MIR. Our experimental machine has 40 Intel® Xeon® E5-2660 cores running at 2.6 Ghz, with 396 GB of RAM running Linux 2.6.32 kernel. We used the Python 2.7.10 interperter to run our experiments.

Datasets: The real-world dataset used in our experiments — called ML-100 — is the 100K example from the MovieLens [19] repository, which contains 100,000 ratings on a scale of [1,5] from 943 users

(scenarios) on 1682 movies (elements) where each user has rated ≥ 20 movies. We binarized this dataset by setting all ratings < 3 to 0, which left us with $82,520$ ratings, where the average user had 87.5 ratings with a standard deviation of 81.2, which suggests a highly-skewed distribution. With this dataset, we use the power-law ($Pr[X = x; \alpha] = \alpha x^{\alpha-1}$) with $\alpha = 1, 2, 3$; note that when $\alpha = 1$, we get a uniform distribution. To get a better understanding of the performance results, we generate multiple permuatations of scenario distributions for the same value of α. For the ODT problem, we also use a synthetic dataset — SYN-K — that is parameterized by k; this is based on a hard instance for the greedy algorithm [23]. Given k, we generate $m = 2k + 3$ sets, $n = k + 2$ elements, with $4k + 4$ non-zeros as follows: (a) elements $i \in [1, k]$ are contained in scenarios $2i - 1$ and $2i$, (b) element $k + 1$ is contained in all odd numbered scenarios, and (c) element $k + 2$ is contained in all even numbered scenarios and scenario $2k + 3$. The probabilities for the scenarios are as follows: $Pr[2i - 1] = Pr[2i] = 2^{-i-2}$ for $i \in [1, k]$, $Pr[2k + 1] = Pr[2k + 2] = 2^{-k-2} - \epsilon$, where $0 < \epsilon < 2^{-k-2}$, and $Pr[2k + 3] = 2^{-1} + 2\epsilon$.

5.1 Optimal Decision Trees

Algorithms. The crux of solving the ODT problem is in chosing an element at each step and making updates to the problem state depending on whether $e \in S_{i^*}$. In our experiments, we compare and contrast the results of 3 algorithms that use different objectives to choose elements: (a) ODT-adsub, which uses the objective described in (3), (b) ODT-greedy, which uses the classic greedy objective [23, 10, 1, 8, 16], and (c) ODT-ml, a "machine learning" algorithm that operates as follows. ODT-ml, which is parameterized by K uses k-Means [2] to *a priori* partition U into K clusters. Each cluster $c_j, j \in [1, K]$ is initially given a weight $w_{c_j} = 1$. To choose the next element $e \in U \setminus E$, a cluster $j \in [1, K]$ is first chosen by sampling non-uniformly according to w_{c_j}. Next, an element $e \in c_j$ is chosen uniformly at random. If $e \in S_{i^*}$, the c_j is rewarded by setting $w_{c_j} = 2w_{c_j}$, else c_j is penalized by setting $w_{c_j} = 0.5w_{c_j}$.

dataset	ODT-adsub	ODT-greedy	ODT-ml
SYN-50	2.75	27.50	21.00
SYN-100	2.75	52.50	21.64
SYN-150	2.75	77.50	56.43
SYN-200	2.75	102.5	53.01
SYN-250	2.75	127.5	69.52

Table 1: Expected cost of running ODT-adsub, ODT-greedy, and ODT-ml algorithms on SYN-K. For ODT-ml, 10 clusters were used, which yielded the best results; various cluster sizes were tried. Results demonstrate that ODT-adsub is superior to its competitors in scenarios that mimic SYN-K.

Results. Table 1 depicts the results for ODT-adsub, ODT-greedy, and ODT-ml on SYN-K datasets for $K = 50, 100, 150, 200, 250$. For this highly-skewed distribution, ODT-adsub outperforms ODT-greedy and ODT-ml (with 10 clusters) by a large margin. It is interesting to note that even the simple ODT-ml algorithm outperforms the well-known ODT-greedy algorithm. Table 2 depicts two sets of results using the ML-100 dataset. The left panel shows the expected cost for the generalized ODT problem with a uniform distribution over scenarios and different thresholds, t_i, on the ML-100 dataset. The first threshold ($t_i = 1$) is the standard setting. For the other two settings, t_is are drawn uniformly at random from the interval indicated in Table. In the left-panel, when the threshold is larger, the expected cost decreases for all three algorithms. Note that although ODT-greedy, which is the best known practical algorithm for ODT, performs the best, ODT-adsub is extremely competitive. Combined with the fact that ODT-greedy performs poorly on worst-case instances (Table 1), we think ODT-adsub is a good alternative in practice. The right panel shows the performance when $t_i = 1$ and the scenarios are drawn from power-law distributions with $\alpha = 2, 3$. For each value of α, 3 random permutations are used to test the stability of the expected cost of each

T	ODT-adsub	ODT-greedy	ODT-ml
1	9.967	9.967	18.219
$[1, 5)$	8.936	8.870	17.121
$[5, 10)$	7.524	7.516	15.466

α	ODT-adsub	ODT-greedy	ODT-ml
$\alpha = 2$	9.841	9.778	18.283
	9.827	9.759	18.359
	9.834	9.775	18.144
$\alpha = 3$	9.651	9.473	18.455
	9.642	9.484	18.486
	9.672	9.497	18.725

Table 2: Expected cost of ODT-adsub, ODT-greedy, and ODT-ml (10 clusters) on ML-100 dataset. Left table shows the performance when scenarios are drawn uniformly at random; the value of T indicates the threshold for stopping. Right table shows the performance of the algorithms when scenarios are drawn from a power law distribution with $\alpha = 2, 3$ ($t_i = 1$ for all i). For each α, we permute scenario probabilities thrice; each permutation assigns a scenario a different probability.

algorithm. From this panel, we can conclude that the performance of ODT-adsub is comparable to that of ODT-greedy for ML-100 dataset when the scenarios are distributed as power law, with ODT-ml placing a distant third. In practive, we could — as we have all the data — fit a custom distribution for the scenarios to get the lowest expected cost.

5.2 Adaptive Multiple Intent Ranking

Algorithms. As for ODT, solving the MIR problem involves carefully choosing an element at each step. In our experiments, we compare and contrast the results of 4 algorithms that use different methods to chose elements: (a) MIR-adsub, which is described in (3), (b) MIR-static, which statically ranks the elements using [4] and choses elements in rank order, (c) MIR-adstatic, which improves on MIR-static by using feedback to eliminate elements from the static list if they belong to invalid scenarios, and (d) MIR-ml, a "machine learning" algorithm that uses the multiplicative scheme described in Section 5.1.

K_i	adsub	static	adstatic	ml
$\lvert S_i\rvert$	92.50	932.42	93.70	101.78
$[\lvert S_i\rvert/2, \lvert S_i\rvert)$	70.38	386.86	71.58	79.70
$[\lvert S_i\rvert/4, \lvert S_i\rvert)$	60.98	304.92	62.19	69.67
$[1, \lvert S_i\rvert/2)$	26.21	84.29	27.31	35.62
$[1, \lvert S_i\rvert/4)$	14.48	38.60	15.54	22.93

α	adsub	static	adstatic	ml
$\alpha = 2$	92.50	932.42	93.70	101.78
	26.21	84.29	27.31	35.62
	14.48	38.60	15.45	22.93
$\alpha = 3$	86.20	905.14	87.78	96.18
	92.72	934.54	94.42	101.52
	91.23	927.03	92.98	100.67

Table 3: Expected cost of MIR-adsub, MIR-static, MIR-adstatic, and MIR-ml (10 clusters) on ML-100 dataset (MIR- prefix dropped in column headers due to space constraints.). Left table shows the performance when scenarios are drawn uniformly at random; the value of K_i indicates the threshold for satisfaction. The right table shows the performance of the algorithms ($K_i = |S_i|$) when scenarios are drawn from a power law distribution with $\alpha = 2, 3$. For each α, we permute scenario probabilities thrice; each permutation assigns a scenario a different probability.

Results. The results of running MIR-adsub, MIR-static, MIR-adstatic, and MIR-ml on the ML-100 dataset are given in the two tables in Table 3. The table in the left-panel shows the expected cost when the scenarios are drawn from a uniform distribution and the satisfaction thresholds, K_is, are varied, with $K_i = |S_i|$ being the classic setting. As expected, when K_i is relaxed, the expected cost decreases for all algorithms. MIR-adsub consistently performs better than all other algorithms in this setting, with MIR-adstatic being a close second. The performance of MIR-static, which is significantly worse than its counterparts demonstrates the importance of *adaptive* algorithms. The table in right panel shows the performance when $K_i = |S_i|$ and the scenarios are drawn from power-law distributions with $\alpha = 2, 3$. For each α, three random permutations are used to test

the stability of the expected cost of each algorithm. The instability of the expected cost across permutations of the user distributions is indicative of the inherent skew in the dataset. Still, MIR-adsub consistently outperforms the other three algorithms.

Acknowledgement: We thank Lisa Hellerstein for a clarification on [15] regarding the OR construction of submodular functions.

References

[1] M. Adler and B. Heeringa. Approximating optimal binary decision trees. *Algorithmica*, 62(3-4):1112–1121, 2012.

[2] D. Arthur and S. Vassilvitskii. k-means++: The advantages of careful seeding. In *SODA*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.

[3] Y. Azar and I. Gamzu. Ranking with submodular valuations. In *SODA*, pages 1070–1079, 2011.

[4] Y. Azar, I. Gamzu, and X. Yin. Multiple intents re-ranking. In *STOC*, pages 669–678, 2009.

[5] N. Bansal, A. Gupta, and R. Krishnaswamy. A constant factor approximation algorithm for generalized min-sum set cover. In *SODA*, pages 1539–1545, 2010.

[6] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012.

[7] G. Bellala, S. K. Bhavnani, and C. Scott. Group-based active query selection for rapid diagnosis in time-critical situations. *IEEE Trans. Information Theory*, 58(1):459–478, 2012.

[8] V. T. Chakaravarthy, V. Pandit, S. Roy, P. Awasthi, and M. K. Mohania. Decision trees for entity identification: Approximation algorithms and hardness results. *ACM Transactions on Algorithms*, 7(2):15, 2011.

[9] F. Cicalese, E. S. Laber, and A. M. Saettler. Diagnosis determination: decision trees optimizing simultaneously worst and expected testing cost. In *ICML*, pages 414–422, 2014.

[10] S. Dasgupta. Analysis of a greedy active learning strategy. In *NIPS*, 2004.

[11] B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008.

[12] U. Feige. A threshold of ln n for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[13] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res. (JAIR)*, 42:427–486, 2011.

[14] D. Golovin, A. Krause, and D. Ray. Near-optimal bayesian active learning with noisy observations. In *NIPS*, pages 766–774, 2010.

[15] N. Grammel, L. Hellerstein, D. Kletenik, and P. Lin. Scenario submodular cover. *CoRR*, abs/1603.03158, 2016.

[16] A. Guillory and J. Bilmes. Average-Case Active Learning with Costs. In *Algorithmic Learning Theory*, pages 141–155. Springer Berlin / Heidelberg, 2009.

[17] A. Guillory and J. Bilmes. Simultaneous learning and covering with adversarial noise. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 369–376, 2011.

[18] A. Gupta, V. Nagarajan, and R. Ravi. Approximation algorithms for optimal decision trees and adaptive tsp problems. In *ICALP (1)*, pages 690–701, 2010.

[19] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 5(4):19, 2015.

[20] L. Hyafil and R. L. Rivest. Constructing optimal binary decision trees is NP-complete. *Information Processing Lett.*, 5(1):15–17, 1976/77.

[21] S. Im, V. Nagarajan, and R. van der Zwaan. Minimum latency submodular cover. *ICALP*, pages 485–497, 2012.

[22] Sh. Javdani, Y. Chen, A. Karbasi, A. Krause, D. Bagnell, and S. S. Srinivasa. Near optimal bayesian active learning for decision making. In *AISTATS*, pages 430–438, 2014.

[23] S. R. Kosaraju, T. M. Przytycka, and R. S. Borgstrom. On an Optimal Split Tree Problem. In *Proceedings of the 6th International Workshop on Algorithms and Data Structures*, pages 157–168, 1999.

[24] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer-Verlag, Berlin, 2003.

[25] M. Skutella and D. P. Williamson. A note on the generalized min-sum set cover problem. *Oper. Res. Lett.*, 39(6):433–436, 2011.

[26] L.A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.